

HP StorageWorks Oracle ZDB Solution with Business Copy EVA 2.3 and VERITAS NetBackup on HP-UX

Product Version: 2

Second Edition (October 2004)

Part Number: 377991-002

This implementation guide provides information regarding HP StorageWorks Business Copy EVA 2.3 snapshot technology with the Oracle 9i database agent for VERITAS NetBackup on HP-UX 11i, using HP StorageWorks Enterprise Virtual Array (EVA) disk arrays. This solution uses scripting to enable the integration of NetBackup and Oracle through Business Copy EVA.

© Copyright 2004 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Compaq Computer Corporation is a wholly-owned subsidiary of Hewlett-Packard Company.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Intel® and Celeron® are U.S. registered trademarks of Intel Corporation.

Intel® and Itanium® are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries and are used under license.

Intel® Itanium™ Processor Family is a trademark in the U.S. and other countries and is used under license.

Microsoft®, MS-DOS®, MS Windows®, Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

Oracle® is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

UNIX® is a registered trademark of The Open Group.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

Oracle ZDB Solution with Business Copy EVA 2.3 and VERITAS NetBackup on HP-UX
Implementation Guide
Second Edition (October 2004)
Part Number: 377991-002

Contents

About this Guide.	5
Intended audience	5
Prerequisites	5
Related documentation	6
Conventions	6
Document conventions	6
Text symbols	6
Equipment symbols	7
Getting help	8
HP technical support	8
HP storage web site	8
HP authorized reseller	8
1 Introduction	9
HP support of zero downtime backup with Oracle	9
Overview of Enterprise Backup Solutions	9
HP HSV Element Manager and Storage Management Appliance overview	10
Business Copy overview	11
History	12
Purpose	12
ZDB topology	13
Planning an Oracle ZDB solution	14
Design consideration factors	14
Snapshots and snapclones	14
Database server prerequisites for the Oracle zero downtime backup solution	15
Important planning considerations	15
2 Implementing an Oracle ZDB Solution	17
Solution components	17
Important Terms	18
Solution execution steps	18
Setting up Business Copy for the HP StorageWorks EVA Disk Array	20
Installing Business Copy 2.3 on the HP HSV Element Manager	20
Install the BC server software using the following procedure:	20
Installing Business Copy 2.3 host agent on the backup and database servers	22

Creating the Business Copy snapshot jobs	25
Identifying the EVA LUNs on the database server	25
Starting the SMA Business Copy web GUI	26
Creating the snapshot jobs on the EVA	26
Verifying the snapshot jobs on the EVA	28
Creating a Business Copy snapshot script on the application server	30
Using the Business Copy CLI Utility <code>evmc1</code> to run and verify snapshot jobs	30
Creating a script to automate run and undo of snapshot jobs	30
Setting up the database and backup servers	34
Preparing the Hardware on the database and backup servers	34
Installing NetBackup components	35
Creating NetBackup policies	35
Creating the ZDB controlling policy	36
Creating the ZDB related data file backup policy	38
Creating the ZDB archive log backup policy	39
Execution and flow of the NetBackup policy pre and post processing scripts	40
Creating the NetBackup policy pre and post processing scripts	41
Creating scripts with the VERITAS NetBackup Oracle Database Agent Script Utility	41
Modifying the ZDB policy preprocessing script	42
Modifying the ZDB policy postprocessing script	47
Running the NetBackup policies to accomplish a ZDB of the Oracle database	53
Running the backup manually	53
Automating the backup	53
Verifying database integrity with snapshot volumes	54
3 Oracle Database Restore and Recovery	55
Database recovery overview	55
Restoring the database files from the snapshot area	55
Restoring the database files from tape	56
Index	57

About This Guide

This implementation guide provides information regarding scripted Oracle zero downtime backup (ZDB) with Business Copy 2.3 and VERITAS NetBackup on HP-UX using HP StorageWorks EVA 5000 disk arrays.

Note: This guide should be used as a *supplement* to the support documentation provided with your solution components.

Intended audience

This guide is intended for use by system administrators implementing an EBS configuration, who are experienced with the following:

- Tape backup technologies and tape libraries
- SAN environments and backup software
- Fibre Channel technology
- HP-UX administration
- HP StorageWorks EVA 5000 configuration

Prerequisites

Before beginning, make sure you have:

- Reviewed the EBS Compatibility Matrix
- Properly installed and configured your EBS hardware per the *HP StorageWorks EBS Design Guide*

Related documentation

In addition to this guide, HP provides corresponding information:

- EBS Compatibility Matrix
- HP blueprints
- *HP StorageWorks EBS Design Guide*
- *HP StorageWorks SAN Design Guide*
- Implementation Guides for supported backup applications
- Installation Guides for EBS hardware components

Conventions

Conventions consist of the following:

- [Document conventions](#)
- [Text symbols](#)
- [Equipment symbols](#)

Document conventions

This document follows the conventions in [Table 1](#).

Table 1: Document conventions

Convention	Element
Blue text: Figure 1	Cross-reference links
Bold	Menu items, buttons, and key, tab, and box names
<i>Italics</i>	Text emphasis and document titles in body text
Monospace font	User input, commands, code, file and directory names, and system responses (output and messages)
<i>Monospace, italic font</i>	Command-line and code variables
Blue underlined sans serif font text (http://www.hp.com)	Web site addresses

Text symbols

The following symbols may be found in the text of this guide. They have the following meanings:



WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or death.



Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

Tip: Text in a tip provides additional help to readers by providing nonessential or optional techniques, procedures, or shortcuts.

Note: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Equipment symbols

The following equipment symbols may be found on hardware for which this guide pertains. They have the following meanings:



Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of personal injury from electrical shock hazards, do not open this enclosure.



Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.



Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of personal injury from a hot component, allow the surface to cool before touching.



Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of personal injury from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.



Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Getting help

If you still have a question after reading this guide, contact an HP authorized service provider or access our web site: <http://www.hp.com>.

HP technical support

Telephone numbers for worldwide technical support are listed on the following HP web site: <http://www.hp.com/support/>. From this web site, select the country of origin.

Note: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

HP storage web site

The HP web site has the latest information on this product, as well as the latest drivers. Access storage at: <http://www.hp.com/country/us/eng/prodserv/storage.html>. From this web site, select the appropriate product or solution.

HP authorized reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-282-6672
- In Canada, call 1-800-863-6594
- Elsewhere, see the HP web site for locations and telephone numbers: <http://www.hp.com>.

Introduction

1

The HP StorageWorks scripted Oracle zero downtime backup (ZDB) solution provides the means to backup an Oracle database with minimal database and database server impact. ZDB is made possible by combining the HP StorageWorks Enterprise Virtual Array (EVA) enterprise disk array snapshot capabilities with the HP StorageWorks Business Copy EVA software utility. This enables the ability to place the Oracle database tablespaces in backup mode briefly while snapshots of the database data files are created. The snapshots can then be presented to a backup server for copying the data to tape. After the backup, the database copy could be re-used as a point-in-time copy if disk space is not an issue. The backup server offloads precious processing resources from the database server, and the Oracle database remains available during the entire process. ZDB solutions are one of many different data protection solutions offered by HP. ZDB solutions are unique in that they tie many software and hardware products together. Choosing a solution based on HP technology provides supportability and known compatibilities.

HP support of zero downtime backup with Oracle

Solutions utilizing zero downtime backup (ZDB) technology consist of hosts, an EVA disk array, Business Copy software, a tape library, and a backup application. HP supports the ZDB Solution through the EVA enterprise disk array and a number of its tape library storage products. This support is made up of SAN interconnects such as fibre channel switches, Host Bus Adapters, disk array controllers and tape controllers. HP EVA disk arrays serve as the hardware snapshot provider, and HP Business Copy software serves as the snapshot initiator in the ZDB based solutions. EBS applications such as VERITAS NetBackup can be used as the backup application, and are supported by EBS.

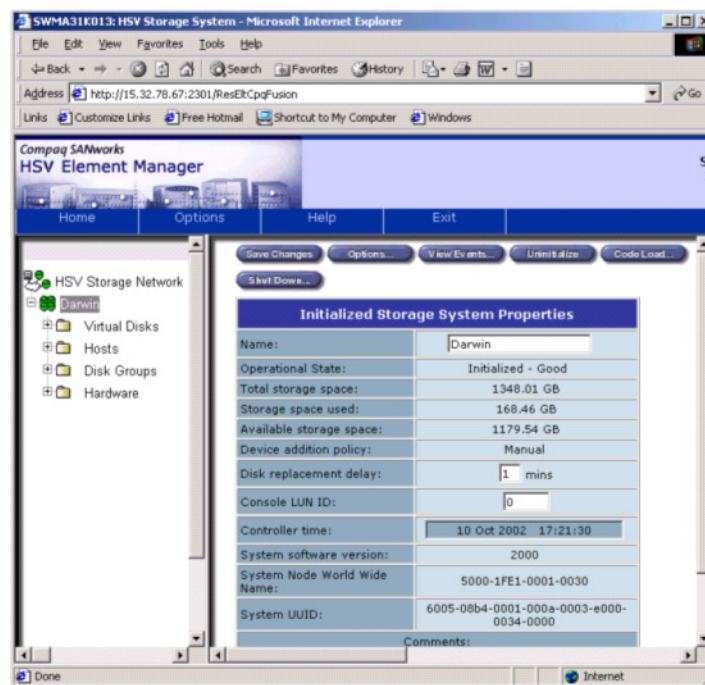
Overview of Enterprise Backup Solutions

ZDB solutions require storage components such as tape libraries, disk arrays and servers, all on a common Fibre Channel storage area network (SAN). HP StorageWorks Enterprise Backup Solution (EBS) is the HP traditional SAN backup solution, where the servers are in the data path from the source to the target. Setting up and configuring a ZDB backup environment is the same as in most EBS environments. Properly setting up a Fibre Channel (FC) SAN backup solution can be challenging. Typically components are purchased at different times and arrive separately, or the components are purchased from different vendors. Each piece of hardware arrives with its own documentation for setup and deployment. These challenges may require additional time and money. HP is committed to keeping these challenges to a minimum by providing the *HP StorageWorks Enterprise Backup Solution Design Guide* and this implementation guide.

HP HSV Element Manager and Storage Management Appliance overview

The HSV Element Manager and Storage Management Appliance provides a web interface to the EVA configuration and monitoring functions. The appliance manages the EVA in band over the SAN. The web interface served by the appliance is accessible over the LAN with a web browser. The left most frame of the Element Manager web GUI contains four folders for each EVA. In this screen shot the EVA called 'Darwin' is shown with the four folders 'Virtual Disks', 'Hosts', 'Disk Groups', and 'Hardware'.

Note: Business Copy 2.3 must be on the same server as the Element Manager, however it is not required to be on the Storage Management Appliance. Refer to the Business Copy administration guide.

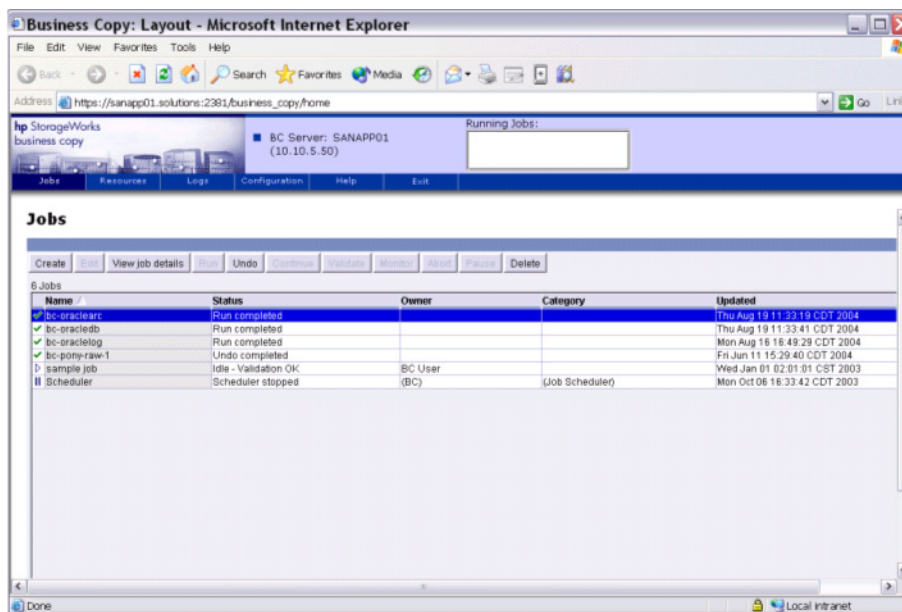


Disk Groups - Physical disks are allocated into pools with common protection levels. These pools are called 'Disk Groups'.

Virtual Disks -The unit of disk space that is seen by hosts on the SAN is called a 'Virtual Disk'. Virtual Disks are simulated disk drives within a physical 'Disk Group'. Properties of a 'Virtual Disk' include the size of the LUN that will be presented to the SAN, The RAID level of the LUN, the world wide LUN number of the LUN, and the hosts that can see it.

Hosts - The 'Hosts' folder contains definitions of the servers that can have 'Virtual Disks' presented to them. A host definition includes the HBA WWN, the network name and IP of the host, and the OS that is running on the host.

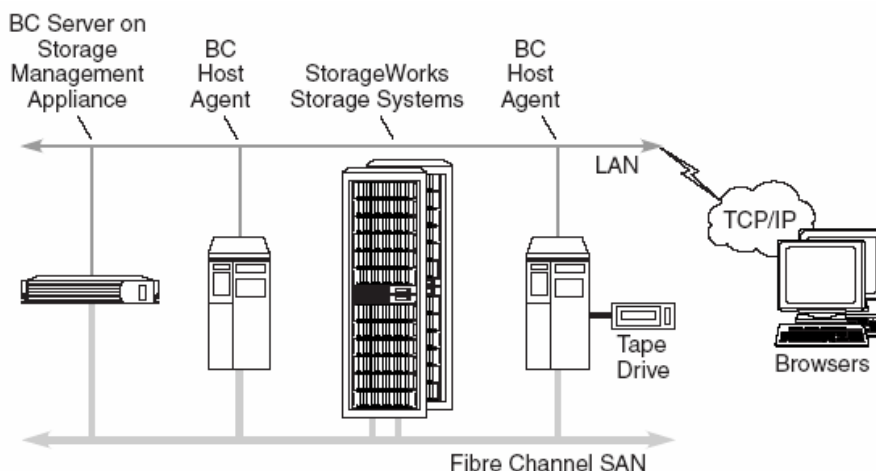
The HSV Element Manager and Storage Management Appliance also provides a web interface for Business Copy configuration and execution. The Business Copy web GUI contains a list of the jobs that have been created. From the GUI new jobs may be created, existing jobs may be edited or deleted, and jobs may be run.



Business Copy overview

Business Copy is an application that makes point-in-time copies of storage volumes.

These copies, called Business Continuance Volumes (BCVs), can be mounted dynamically on any other supported host in the BC storage network. To replicate storage volumes, you must create a BC network. The following figure shows a typical BC network.



The BC network consists of the following hardware and software components:

- One Storage Management Appliance (SMA)
- BC server software, which runs on the SMA

- Device managers/element managers, which run on the SMA
- One or more host computers
- BC host agent software, which run on host computer systems
- One or more HP StorageWorks storage systems

History

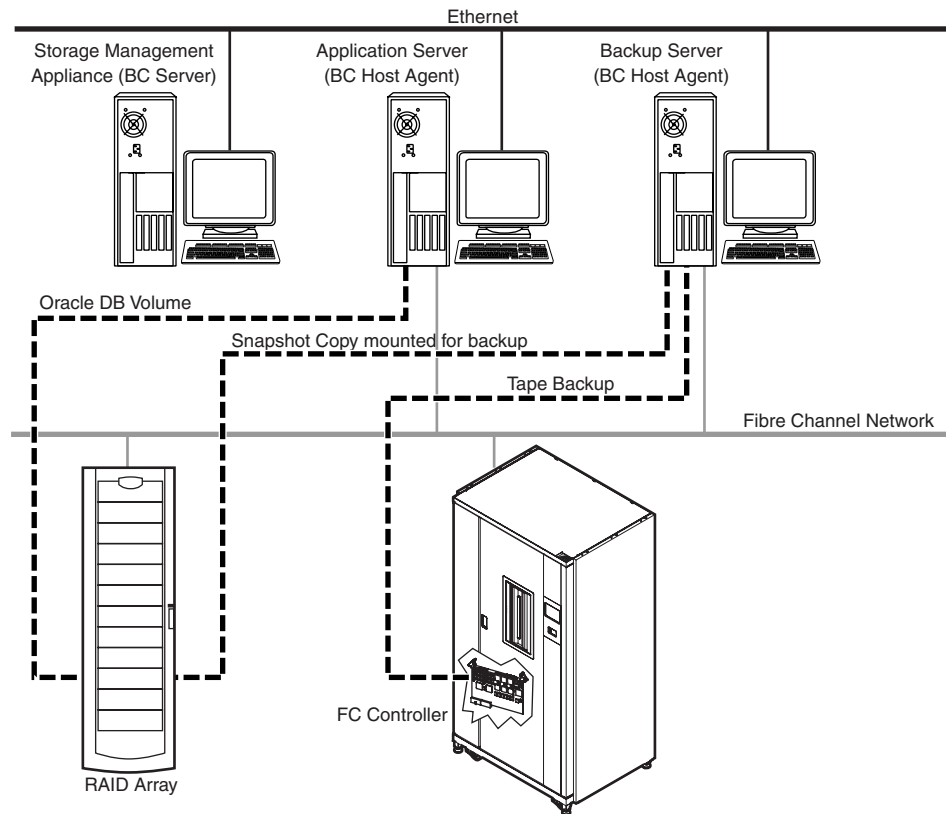
HP engineering teams have developed a comprehensive approach to ensuring that all hardware, firmware, and software components are properly fitted into an Enterprise Backup Solution (EBS). Teams test the supported configurations and develop Best Practices to follow when setting up your own EBS. The teams also test backup solution software and provide best practices to ensure that your EBS runs smoothly.

Purpose

This guide is intended to address many of the integration issues that you may encounter when setting up your EBS and to provide suggestions for the best solution. This guide does not provide specific documentation for installing and configuring your data protection software or tape library hardware. You will be referred to the appropriate documentation when necessary. The guide covers areas where special configuration procedures that might not be covered in the vendor documentation will help in your goal of setting up an efficient EBS.

Refer to the *HP StorageWorks Command View EVA User Guide* for proper disk array setup and configuration.

ZDB topology



Planning an Oracle ZDB solution

There are several factors to consider when designing a ZDB solution. These factors include database server impact, cost, backup and restore performance, array integration, and ease of configuration and use. While all of these factors are important, because of the complexity of these solutions they may be easy to use, but can be difficult to setup.

Design consideration factors

While the cost of the solution is always a factor, it is database server impact and database availability that seem to take center stage in ZDB solutions today. It is easy to see how the database server is removed from the heavy demands of backups in these solutions as the demand is handled by an off-host server; however database availability is the most noted benefit of ZDB solutions. Good performance can be achieved by using a capable backup server or multiple backup servers.

Restores are obviously a very important part of the solution's equation but are sometimes overlooked due to the benefits of the backups. Restores may be accomplished by retrieving data from tape directly to the database server, or, if the database snapshot still exists data can be restored directly from the snapshot area to the database server.

Finally, array integration can make a big difference in how the ZDB solution behaves. ZDB solutions require a hardware snapshot of the primary data volume. You must consider whether to use snapshots or snapclones. A snapshot is a nearly instantaneous virtual copy of a disk or LUN. A snapshot does not actually copy an entire disk, but rather makes a copy of the meta data for that LUN. When blocks are changed (updated) on the original disk, the modified blocks are first written to the snapshot area and then are flagged in the meta data to denote that the original blocks are now in the snap area and not on the original disk.

Snapshots and snapclones

Snapshots may be either space efficient or fully allocated. A fully allocated snapshot pre-allocates all possible required space when the snapshot is created. A space efficient snapshot only allocates space on an as-needed basis when data is modified. Zero gigabytes are allocated initially. If 2 GB of data are changed, then 2 GB are allocated in the snap area.

A snapclone is a combination of a snapshot and a clone. A snapclone starts out as a fully allocated snapshot, allowing it to be presented virtually instantaneously without needing any time to synchronize. The cloning operation takes place in the background. At some point in the near future, the copy operation completes and all data now resides on separate volumes as with a traditional clone.

There is an important difference between Virtually Instantaneous Snapclone and the traditional clone. With traditional, the clone copy is not available until the copy is complete. With Virtually Instantaneous Snapclone, the snapclone data can be accessed virtually immediately, with full redundancy available when the background copy completes.

Performance of the primary data volumes may be reduced during backup of data from the snapshot area to tape. This is because the backup may access the primary data volume to retrieve some of the data required for backup.

See HP StorageWorks Enterprise Virtual Array documentation for more information on snapshots and snapclones.

Database server prerequisites for the Oracle zero downtime backup solution

- The Oracle database must be running in archivelog mode.
- For file system backups:
 - The Oracle volumes must reside on VxFS file systems.
 - If Online JFS is installed on the database server, the Oracle volume VxFS file systems should be mounted using the VERITAS direct I/O options:
 - mincache=direct
 - convosync=direct
 - The Oracle volume VxFS file systems should use control intent logging which is enabled by mounting the file system with the *log* or *delaylog* option.

Note: To eliminate the need to unmount the file-system prior to back up, and to guarantee file-system and data integrity, HP recommends using Online JFS and "direct I/O" with the *log* option. There are cases where the alternative direct I/O option *delaylog* will not work, such as a crash during file/directory creation or deletion.

Important planning considerations

There are seven high level steps that must happen for Oracle zero downtime backups to take place. They are as follows:

1. The database tablespaces are placed into backup mode.
2. A snapshot of the primary data volumes is created and presented to the backup server.
3. The database tablespaces are put in normal mode.
4. The snapshot data is copied to tape.
5. A database checkpoint is taken and the current logs are switched.
6. A snapshot of the archive logs is created and presented to the backup server.
7. The snapshot archive logs are copied to tape.

Care must be taken while planning the ZDB solution to ensure the components of the solution can accomplish these seven steps.

Special attention must be focused on presenting the correct targets to the appropriate initiators when using zoning on your SAN. The EVA disk array and tape library interface controllers must be presented to both the database server and the backup server. For ZDB to function properly, the data to be backed up must be located on an EVA disk array that is part of the SAN environment. This ensures that the snapshot data can be copied to the tape library by the backup server, and can be retrieved from the tape library by the database server or the backup server.

Implementing an Oracle ZDB Solution

2

Note: The configuration rules and recommendations are made based on the solution integration testing conducted by HP. Certain limitations apply and are noted where applicable. This implementation guide can be leveraged as a template for similar solutions where there may be minor changes in design, such as the operating system version or HBAs used.

Solution components

- HP StorageWorks Business Copy EVA 2.3
- HP StorageWorks Enterprise Virtual Array (EVA) 5000 disk array
 - VCS 3.020
- HP OpenView Storage Management Appliance Version 2.1
- HP 9000 PA-RISC Database Application and Backup Servers
 - HP-UX 11i for PA-RISC (11.11)

Note: VERITAS NetBackup does not currently support HP-UX 11.23 on IA-64 as a media server.

- Logical Volume Manager (LVM)
 - Online JFS or JFS
 - VxFS filesystems
- Oracle Database 9i
- VERITAS NetBackup 5.1
 - Media Server
 - Shared Storage Option
 - Oracle Database Agent
- HP StorageWorks ESL9000 Series LTO tape libraries with the HP StorageWorks e2400-160 Fibre Channel Interface Controller and the HP StorageWorks Interface Manager card
- A6795A and/or A6826A FC host bus adapters

Important Terms

- **Database server**—Oracle database server.
- **Backup Server**—The offhost that does data backups from the snapshot area to tape.
- **Virtual Disk**—A simulated disk that is created from physical disks in an EVA disk array for use by a host.
- **Business Continuance Volume (BCV)**—An EVA disk array virtual disk that is created from a production volume and can subsequently be used for tasks such as backups.
- **Snapshot**—A StorageWorks term meaning a nearly instantaneous controller-based method of creating a virtual copy of a virtual disk.
- **Snapclone**—A StorageWorks term meaning a virtual copy of a virtual disk that begins as a fully allocated snapshot then becomes an independent virtual disk.
- **BC Job**—A file created by Business Copy (BC) that represents a user request to perform a task or a series of tasks. For example, a BC job named snapclone-oracledb might be used to create a snapclone of the Oracle database volumes and then mount the snapclone on a backup server.
- **BC Run**—The process of executing the steps of a previously created BC job.
- **BC Undo**—The mode of running a BC job for the purpose of “undoing” steps that have been completed.
- **NetBackup preprocessing script**—A script that is run by a NetBackup policy prior to performing backups.
- **NetBackup postprocessing script**—A script that is run by a NetBackup policy after the completion of backups.

Solution execution steps

This solution executes the following steps:

1. **If previous database and archive log snapshots exist, undo the snapshots**—Initiated by the preprocessing script of the database server NetBackup Policy.
2. **Place the database tablespaces in backup mode**—Initiated by the preprocessing script of the database server NetBackup Policy.
3. **Create a snapshot of the database data volumes, and mount the snapshot volumes on the backup server**—Initiated by the preprocessing script of the database server NetBackup Policy.
4. **Copy the database data snapshot to tape**—Initiated by the backup server NetBackup policy which cannot start until the database server NetBackup policy preprocessing script exits.
5. **Copy the database parameter and password files to tape**—Initiated by the database server NetBackup policy after the preprocessing script exits.
6. **Take the database tablespaces out of backup mode, do a database checkpoint and switch the database log file**—Initiated by the postprocessing script of the database server NetBackup policy.
7. **Create a snapshot of the archive logs and mount the archive logs snapshot on the backup server**—Initiated by the postprocessing script of the database server NetBackup policy.
8. **Copy the archive logs snapshot to tape**—Initiated by the NetBackup archive log backup policy on the backup server. This policy is started by the database server NetBackup policy postprocessing script prior to exiting.

9. **Leave the database volume and archive snapshots until the next database backup**—After the backup to tape is completed the snapshots may be undone, if desired. This removes the snapshot copy of the data and frees the disk space used by the snapshots. If the snapshots are left intact until the next backup, the snapshot copy of the data is available for the database recovery, if needed.

Note: Steps 5 through 8 take place in parallel to the backup of the database data files to tape, which is initiated in step 4.

Setting up Business Copy for the HP StorageWorks EVA Disk Array

This solution does not go into the detail of adding hosts to the EVA disk array and creating virtual disks on the EVA disk array. It is assumed that the database server and backup servers have already been added as hosts on the disk array and that the database volume virtual disks have been created and presented to the database server.

For detailed information on setting up the EVA disk array please refer to the *HP StorageWorks Enterprise Virtual Array User Guide*, available from the HP web site:

<http://www.hp.com/go/storage>

Click on the software link under storage products.

Installing Business Copy 2.3 on the HP HSV Element Manager

Note: This is an overview of the Business Copy 2.3 installation. Refer to the *HP StorageWorks Business Copy EVA/MA/EMA Server Installation Guide* for detailed installation information.

This section describes how to install HP storage management applications on the Element Manager, or on the SMA (if the Element Manager resides on the SMA). One or more of the following may be required to complete this procedure:

- Product name—Business Copy Server v2.3
- Product CD-ROM
- Business Copy Replication License Key
- FTP server name and path, if using the Network or FTP server installation method
- Network package file name (SWP), if using the FTP server or Appliance installation method—`bc_23_server.swp`

Note: In some cases, the display may perform an automatic reboot following a successful installation.

Install the BC server software using the following procedure:

1. Close active applications.
 - a. Close all browser windows, Microsoft Management Console (MMC) sessions, Terminal Services sessions, and Java™ applets that are open to, and on, the SMA.
 - b. For new installations only, insert the CD-ROM from the BC product kit into the SMA CD drive, and then go to step 2.
 - c. For update installations only, go to step d.
 - d. Using the SMA Web GUI, stop the BC server service by selecting **Home > Settings > Manage Tools > Business Copy**.
 - e. Click **Start** or **Stop**.
2. From a client computer, launch a Web browser and browse to the SMA using the following format: `http://<MyAppliance_name or IP Address>`
3. Log in to the SMA. If necessary, refer to the BC Network Administration Guide.

4. Click **Settings**.
5. Click **Maintenance > Install Software**.

Note: If reinstalling or updating BC server software, make sure the application is not running before proceeding. Click the Manage Tools link and stop the application service, if necessary.

6. Click **Next** at the bottom of the page.
7. Select the installation option, and then click **Next**.

Note: Only the options presented in this procedure are supported for this product release.

■ **CD-ROM:**

- a. Click **Next** at the bottom of the page.
- b. Select **BC 23 Server** from the drop-down list.

Continue with step 8.

■ **FTP Server:**

- a. Type, or accept, the following information to establish an FTP network connection with the SMA. All entries are case-sensitive.
 - For **FTP server name**, type the fully qualified domain name of the FTP server on which the `bc_23_server.swp` file is located (or enter the IP address of the server).
 - For **Full File Path**, type the default path of the download folder and the `bc_23_server.swp` file name in the text box.
For example: `/<download_folder_path>/bc_23_server.swp`
 - For **User name**, accept the default of anonymous.

If the FTP server does not accept anonymous access, a user name and password must be included. Type the user name and password in the respective fields.

- b. Click **Next** at the bottom of the page.
- c. Select **BC 23 Server** from the drop-down list.

Continue with step 8.

■ **Local Disk:**

- a. Select **BC 23 Server** from the drop-down list.

Note: If the BC 23 Server name does not display in the drop-down list, verify that the `bc_23_server.swp` file is located in the correct directory.

Continue with step 8.

8. Click **Next** to initiate the installation.

The display states:

Installation is in progress.

Time to completion depends on the size of the application and the network connection speed. After approximately 2 to 3 minutes, the display states:

Installation is complete.

StorageWorks Business Copy - Installation Complete.

Note: In some cases, the display may indicate `Rebooting this Appliance`. If an installation fails for any reason, the SMA does not permit a subsequent installation attempt for 1 hour following the failed attempt. Beginning an installation attempt during this 1-hour period displays an error message.

9. Click **Finish**.

10. Does a **business copy** entry exist on the **Tools** page?

- Yes. The BC server software is now installed.
- No. Return to step 4 to repeat the installation or refer to the BC Network Administration Guide for resolving issues.

Installing Business Copy 2.3 host agent on the backup and database servers

Note: This section is an overview of the BC 2.3 host agent installation. Refer to the *HP StorageWorks Business Copy EVA/MA/EMA Host Agent v2.3 for HP-UX Installation Guide* for details on installing the Business Copy 2.3 host agent.

Complete the following procedure to install a BC host agent on a host.

Note: During the installation procedure, press **Enter** to choose the default response. To return to a previous step, enter `back`. To exit the installation script at any time, enter `quit`. Also press **Enter** as necessary to scroll through lengthy text screens.

1. Enter `# sh ./bc_23_hpux_ha_install.bin` to start the installation script. General information displays as the installation script unpacks, followed by BC-specific information.

Note: Unpacking of the Java™ Runtime Environment (JRE) segment can take three to five minutes.

2. Press **Enter** to continue.
3. Review the BC license agreement.
4. Enter `y` or `yes` to agree to these terms and continue the installation; enter `n` or `no` to exit the installation script. Entering `yes` provides a list of BC operational prerequisites.

Note: HP recommends exiting the installation if BC prerequisites have not been met, or if doubts exist as to whether the requirements are met. Take the necessary action to make sure that these prerequisites are met and then restart the installation.



Caution: Failure to properly address BC preinstallation requirements and compatibility considerations can lead to reduced operational capability and failure of BC jobs.

5. Carefully review the BC prerequisites and make sure each item has been addressed before continuing the BC host agent installation.

Note: If any of the operational considerations have not been addressed, exit the installation, address the items, and then restart the installation.

6. Press **Enter** to continue; enter `quit` to exit the installation script. An option to select the BC installation folder displays.
 - For initial installations, `/opt` displays as the BC default installation folder.
 - For reinstallations, the folder used during the previous installation displays as the default installation folder.



Caution: HP does not recommend changing the installation folder when reinstalling the BC host agent. Changing the folder during a BC host agent reinstallation might cause BC to function improperly.

Note: If updating to BC Host Agent v2.3 software, the update will automatically be placed in the folder used during the initial BC installation.

7. Press **Enter** to continue using the BC default installation folder, or enter an absolute path to the desired installation folder; or enter `quit` to exit the installation script.
8. Enter `yes` if the installation folder is correct; enter `no` to reenter the installation folder. After the installation folder is determined, the BC host agent software installation begins. Initially, a progress bar is displayed, followed by an SMA name question.

Note: For reinstallations, the previous SMA name displays as the default.

9. Enter the fully qualified name, qualified name, or IP address of the SMA that supports this BC host agent. A check is performed to verify communication with the SMA name or IP address.

The check provides the opportunity to continue the BC host agent installation, in the event that the LAN is down or the SMA is offline. This check allows an incorrect SMA name to be corrected.

- Entering a 1 or pressing **Enter** continues the installation using the SMA name entered in step 9.
 - Entering a 2 continues the installation at step 9.
10. Press **Enter** to continue; enter `quit` to exit the installation. The BC host agent software installation takes place.

11. Press **Enter** to continue.

The BC host agent is now installed. Refer to the BC Host Agent Release Notes.

Note: The installation script automatically starts the BC daemons. To manually stop and start the BC daemons, refer to the BC Network Administration Guide for details.

12. Press **Enter** to complete the installation.

Creating the Business Copy snapshot jobs

This section describes how to identify the EVA LUNs on the database server, start the SMA Business Copy web GUI, and how to create and verify Business Copy snapshot jobs.

Identifying the EVA LUNs on the database server

In this example the filesystems mounted on `/oracledb`, `/oraclelog` and `/oraclearc` will be the source of the snapshots. The output of `bdf` shows the logical volumes as `/dev/vg_oracledb/lvol1`, `/dev/vg_oraclelog/lvol1` and `/dev/vg_oraclearc/lvol1`.

```
# bdf
```

Filesystem	kbytes	used	avail	%used	Mounted on
/dev/vg00/lvol3	2097152	96368	1985176	5%	/
/dev/vg00/lvol1	1014648	60952	852224	7%	/stand
/dev/vg00/lvol8	4194304	505896	3660200	12%	/var
/dev/vg00/lvol7	4194304	1619048	2555200	39%	/usr
/dev/vg00/lvol6	1048576	190528	851424	18%	/tmp
/dev/vg_oraclelog/lvol1	10481664	34158	10121028	0%	/oraclelog
/dev/vg_oracledb/lvol1	20967424	1846724	18821948	9%	/oracledb
/dev/vg_oraclearc/lvol1	41934848	281216	41328280	1%	/oraclearc
/dev/vg_oracle/lvol1	10485760	4826394	5482738	47%	/oracle
/dev/vg00/lvol5	4194304	1627688	2547216	39%	/opt
/dev/vg00/lvol4	4194304	120768	4041728	3%	/home

Running `lvdisplay -v` against the `/oracledb` logical volume displays the *sdisk* device file bound to the Virtual Disk.

```
# lvdisplay -v /dev/vg_oracledb/lvol1
```

```
--- Logical volumes ---
```

LV Name	/dev/vg_oracledb/lvol1	
VG Name	/dev/vg_oracledb	
... ..		

```
--- Distribution of logical volume ---
```

PV Name	LE on PV	PE on PV
/dev/dsk/c28t0d1	5119	5119
... ..		

Running `lssf` against the device file yields the hardware path.

```
# lssf /dev/dsk/c28t0d1
```

```
sdisk card instance 28 SCSI target 0 SCSI LUN 1 section 0 at address 0/1/0/0.2.59.0.0.0.1 /dev/dsk/c28t0d1
```

The result of running `ioscan` on the hardware path should show a description field that lists the device as a COMPAQ HSV110 and the hardware path should include the SCSI target ID and LUN. If any part of the filesystem is on a non COMPAQ HSV110 then this snapshot solution will not work.

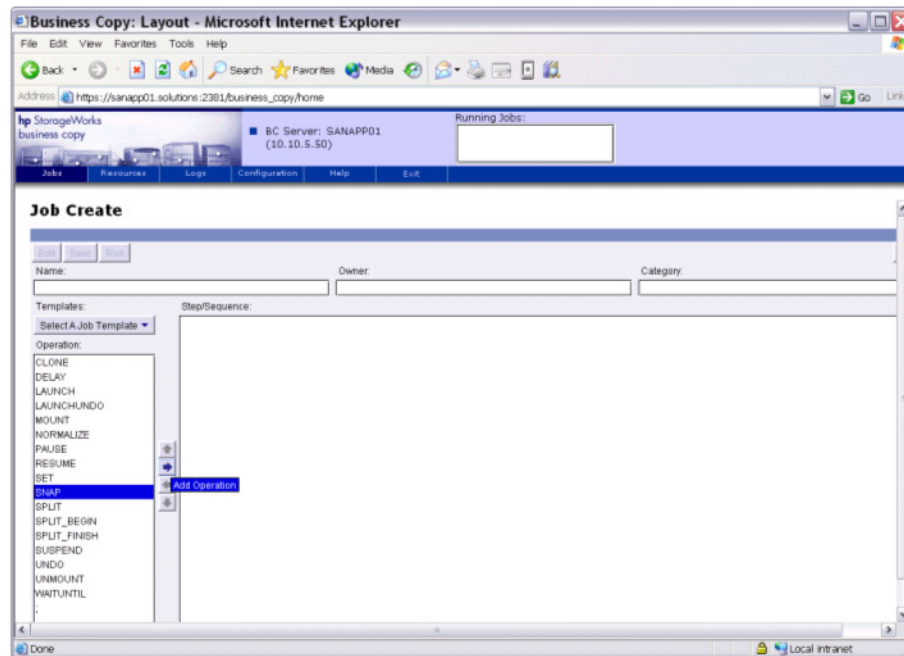
```
# ioscan -fnkH 0/1/0/0.2.59.0.0.0.1
Class I H/W Path Driver S/W State H/W Type Description
=====
disk 99 0/1/0/0.2.59.0.0.0.1 sdisk CLAIMED DEVICE COMPAQ HSV110 (C)COMPAQ
/dev/dsk/c28t0d1 /dev/rdisk/c28t0d1
```

Starting the SMA Business Copy web GUI

1. From a client computer, launch a Web browser.
2. Browse to the SMA, using the following format: `http://<MyAppliance_name>`
3. Log into the SMA.
4. Click **Tools**.
5. Click **business copy**.

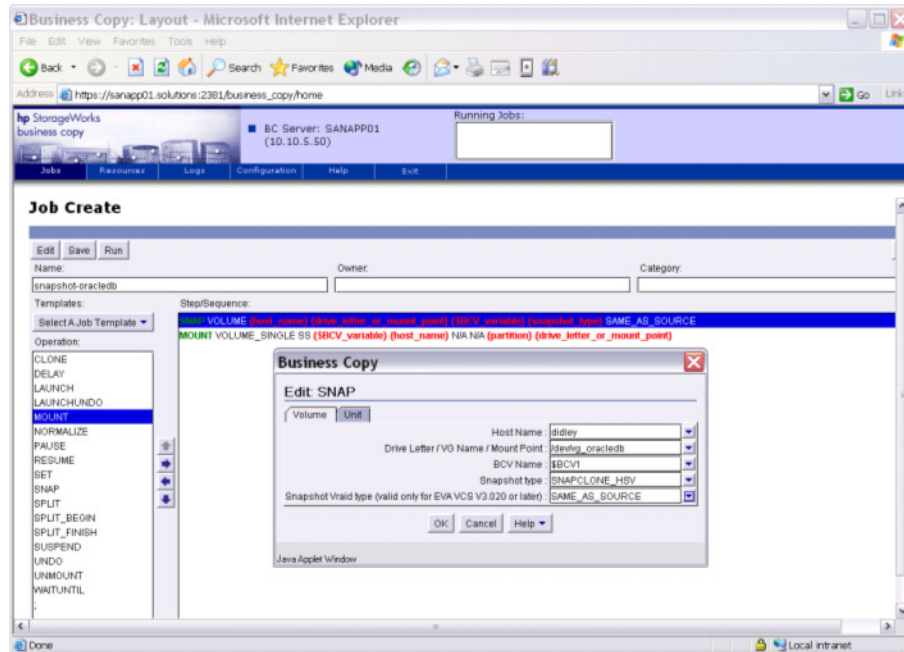
Creating the snapshot jobs on the EVA

1. Click **Create** in the Business Copy Web GUI to go to the *Job Create* window.
2. In the *Job Create* window, give the job a unique name (required) and assign it an owner (optional). It is good practice to include the Oracle SID as a part of the unique name.
3. Highlight the **SNAP** operation and click on the *Add Operation* arrow to add the first step of the job as a snapshot

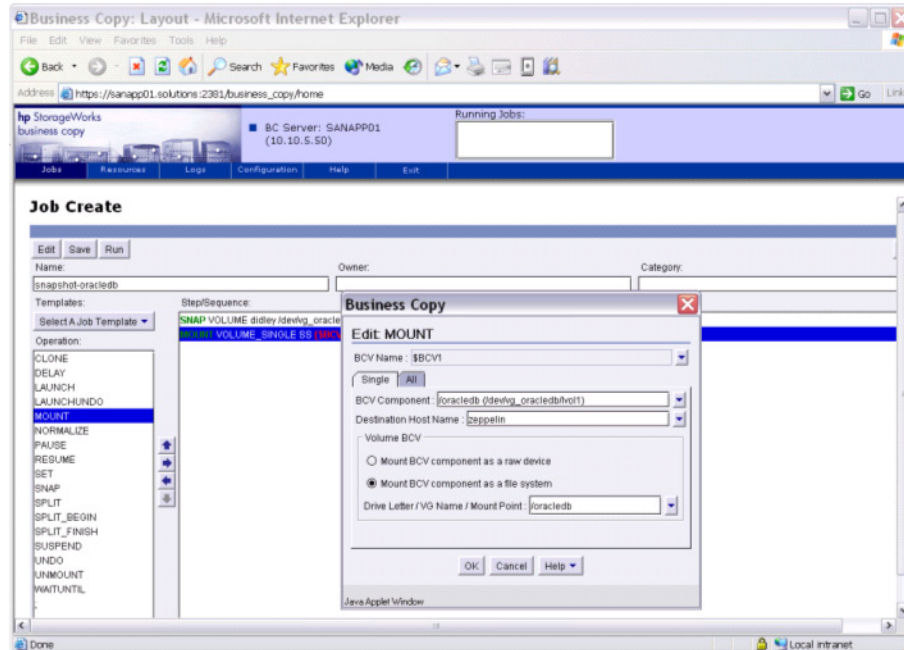


4. Highlight the **MOUNT** operation and click on the *Add Operation* arrow to add the second step of the job as a mount.
5. Double click on the **SNAP** step and add the job specific information.
 - a. **Host Name**—The name of the database server.
 - b. **VG Name**—The path of the LVM volume group (`/dev/vg_oracledb` in this example).

- c. **BCV Name**—The name of the BCV. Select a name from the drop down list of available names.
- d. **Snapshot type**—Fully Allocated (a snapshot in which the disk space is fully allocated when the snapshot is created), Demand Allocated (a snapshot in which disk space is allocated when needed), or Snapclone.
- e. **Snapshot Vraid type**—Must be SAME_AS_SOURCE unless the EVA is running VCS V3.020 or later.



- 6. Double click on the **MOUNT** step and add the job specific information.
 - a. **BCV Name**—Same as the BCV name selected for the SNAP step.
 - b. **BCV Component**—Mount point of the volume on the database server.
 - c. Mount BCV component as a file system.
 - d. **Mount Point**—The mount point of the snapshot on the backup server. This can be the same as the mount point on the database server to make database restores simpler but does not have to be the same.



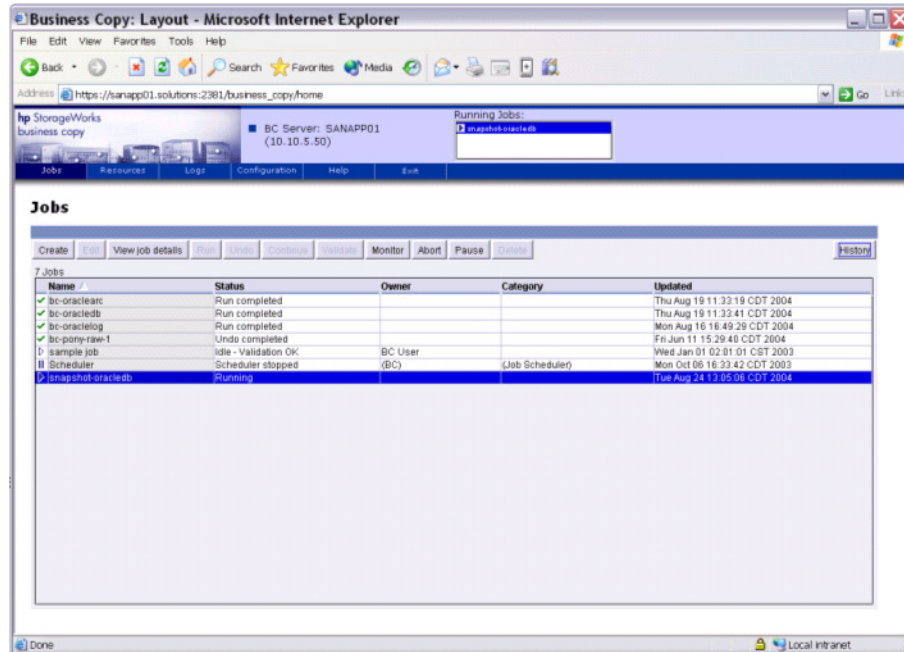
7. Save the job, and click **Jobs** to return to the main Business Copy Web GUI.
8. Repeat job creation steps 1 through 7 for each snapshot required for your database data file and archive log volumes.

Verifying the snapshot jobs on the EVA

1. In the Business Copy Web GUI, highlight the job to verify and click **Validate**. If the validation succeeds the job status will change to *Idle—Validation OK*. If the validation fails the job will need to be edited to fix the problem.

For help debugging the cause of job validation failure, view the *job_jobname_#.txt* log by selecting **Logs** in the Business Copy web GUI.

2. If the job validation succeeded, the next verification step is to run the job. Highlight the job and click **Run**. The job status will change to *Running*. The run will take a few minutes, and if it succeeds the job status will change to *Run Completed*.



3. After the run has completed, `bdf` on the backup server should show the mounted snapshot.

```
# bdf
```

Filesystem	kbytes	used	avail	%used	Mounted on
/dev/vg00/lvol3	204800	100504	103528	49%	/
/dev/vg00/lvol1	298928	71304	197728	27%	/stand
/dev/vg00/lvol8	4710400	975080	3706168	21%	/var
/dev/vg00/lvol7	4194304	1734024	2441096	42%	/usr
/dev/vg00/lvol4	204800	99736	104304	49%	/tmp
/dev/vg_oracle/lvol1	10485760	4825882	5483070	47%	/oracle
/dev/vg00/lvol6	4194304	2198152	1982104	53%	/opt
/dev/vg_oraclelog_BCV_0/lvol1	10481664	34158	10121028	0%	/oraclelog
/dev/vg_oraclearc_BCV_0/lvol1	41934848	281216	41328280	1%	/oraclearc
/dev/vg_oracledb_BCV_0/lvol1	20967424	1846724	18821948	9%	/oracledb

4. Highlight the job in the Business Copy Web GUI and click **Undo**. The job status will change to *Undoing*. The undo will take a few minutes, and if it succeeds the snapshot volume will be unmounted on the backup server and the job status will change to *Undo completed*.
5. Repeat verification steps 1 through 4 for all jobs created.

Creating a Business Copy snapshot script on the application server

To automate the run and undo of snapshots on the EVA disk array from the backup server a script should be created and called by a NetBackup policy. The installation of the Business Copy Host Agent on the backup server included CLI utilities that can be used by a script for getting job status, running jobs and undoing jobs. The CLI commands should be installed in */opt/CPQevm/bin*.

Using the Business Copy CLI Utility `evmcl` to run and verify snapshot jobs

The command used to get job status, run and undo jobs is `evmcl`:

- To return job status: `evmcl appliance_name status job_name`
- To run a job: `evmcl appliance_name execute job_name`
- To undo a job: `evmcl appliance_name undo job_name`
- To list existing jobs: `evmcl appliance_name getjoblist`

Each of the above commands should be run to verify the `evmcl` utility can get job status, run and undo jobs. In the following examples the appliance name is `sanapp01` and the job name is `snapshot-oracledb`. In addition to using `evmcl` to get the job status, you can use the BC Web GUI to see job status and job execution.

Example 1—Run a BC job:

```
# /opt/CPQevm/bin/evmcl sanapp01 status snapshot-oracledb
— Job status should be Undo Completed.

# /opt/CPQevm/bin/evmcl sanapp01 execute snapshot-oracledb
— This should take a few minutes to complete and the job status should be Running.

# /opt/CPQevm/bin/evmcl sanapp01 status snapshot-oracledb
— Job status should be Run Completed and the snapshot should be mounted on the
  backup server.
```

Example 2—Undo a BC job:

```
# /opt/CPQevm/bin/evmcl sanapp01 status snapshot-oracledb
— Job status should be Run Completed.

# /opt/CPQevm/bin/evmcl sanapp01 undo snapshot-oracledb
— This should take a few minutes to complete and the job status should be Undoing.

# /opt/CPQevm/bin/evmcl sanapp01 status snapshot-oracledb
— Job status should be Undo Completed and the snapshot should be unmounted from the
  backup server.
```

Creating a script to automate run and undo of snapshot jobs

The snapshot script should be able to do the following:

- Get the status of one or more BC jobs.
- Run one or more BC jobs (if the status of a job is *Run Completed* the script should undo the job prior to running it).
- Undo one or more BC jobs.

The following is a sample Korn Shell script:

```
#!/bin/ksh

get_job_status() {
    StatJob=$1
    Status=`/opt/CPQevm/bin/evmcl $ApplianceName status $StatJob | /sbin/awk '{print
$11 $12}'`
}

status_check() {
    CheckType=$1
    shift
    CheckJobs=$*
    AllJobsStatus="false"
    while [[ "$AllJobsStatus" = "false" ]]; do
        print "Job status check:"
        AllJobsStatus="true"
        for checkjob in $CheckJobs; do
            get_job_status $checkjob
            checkjobstatus=$Status
            print "\tStatus of $checkjob: $checkjobstatus"
            if [[ "$checkjobstatus" = "Failed" ]]; then
                print "Job $checkjob failed!"
                exit
            elif [[ "$checkjobstatus" != "$CheckType" ]]; then
                AllJobsStatus="false"
            fi
        done
        if [[ "$AllJobsStatus" = "false" ]]; then
            /bin/sleep 20
        fi
    done
}

jobs_undo() {
    UndoJobs=$*
    sleeptime=10
    print "Undoing jobs $UndoJobs."
    for undojob in $UndoJobs; do
        /opt/CPQevm/bin/evmcl $ApplianceName undo $undojob /i
        /bin/sleep $sleeptime
        let sleeptime=$sleeptime+5
    done
    status_check "Undocompleted" $UndoJobs
    print "All jobs successfully completed undo."
}

jobs_run() {
    RunJobs=$*
    print "Running jobs $RunJobs."
}
```

```
for runjob in $RunJobs; do
    /opt/CPQevm/bin/evmcl $ApplianceName execute $runjob /i
    /bin/sleep 10
done
status_check "Runcompleted" $RunJobs
print "All jobs successfully completed run."
}

Usage="Usage: $0 [-d|-u] [-a <appliancename>] <job1> [<job2> ...]
-d = Do snaps (default)
-u = Undo snaps (-d and -u are mutually exclusive)
-s = Do a status check only (-d, -u and -s are mutually exclusive)
-a = Name of SAN appliance (default: sanapp01)"

if [[ $# -lt 1 ]]; then
    print "$Usage"
    exit
fi

ApplianceName=sanapp01
JobType=do

SnapFlag=0
while getopts dusa:h c; do
    case $c in
        d) if [[ $SnapFlag -eq 0 ]]; then
            JobType=do
            SnapFlag=1
        else
            print "Error: -d, -u and -s flags are mutually exclusive!"
            print "$Usage"
            exit 0
        fi;;
        u) if [[ $SnapFlag -eq 0 ]]; then
            JobType=undo
            SnapFlag=1
        else
            print "Error: -d, -u and -s flags are mutually exclusive!"
            print "$Usage"
            exit 0
        fi;;
        s) if [[ $SnapFlag -eq 0 ]]; then
            JobType=stat
            SnapFlag=1
        else
            print "Error: -d, -u and -s flags are mutually exclusive!"
            print "$Usage"
            exit 0
        fi;;
        a) ApplianceName=$OPTARG;;
    esac
done
```

```

        h) print "$Usage"
        exit 0;;
        \?) print "$Usage"
        exit 0;;
    esac
done
shift `expr $OPTIND - 1`

AllJobs=$*

DoneJobs=""
UndoneJobs=""
for job in $AllJobs; do
    get_job_status $job
    JobStatus=$Status
    if [[ "$JobStatus" = "Undocompleted" ]]; then
        UndoneJobs="$UndoneJobs $job"
    elif [[ "$JobStatus" = "Runcompleted" ]]; then
        DoneJobs="$DoneJobs $job"
    fi

    if [[ "$JobType" = "stat" ]]; then
        print "Status of $job: $JobStatus"
    fi
done

if [[ -z "$DoneJobs" && -z "$UndoneJobs" ]]; then
    print "No jobs were found!"
else
    if [[ -n "$DoneJobs" && "$JobType" != "stat" ]]; then
        jobs_undo $DoneJobs
        if [[ "$JobType" = "do" ]]; then
            /bin/sleep 10
        fi
    fi

    if [[ "$JobType" = "do" ]]; then
        jobs_run $AllJobs
    fi
fi

```

Setting up the database and backup servers

The setup of the database and backup servers includes preparing the tape library, and installing and configuring NetBackup. This document assumes that Oracle has been previously installed, that a database instance is up and running, and that the database data file(s) and archive log volumes are located on an EVA disk array.

Preparing the Hardware on the database and backup servers

A shared tape library is part of the Oracle ZDB solution.

Note: This document does not go into the details of installing a tape library on the SAN. Please refer to the HP StorageWorks Enterprise Backup Solution Design Guide for detailed information on installing a tape library.

The tape library robot and tape devices should be presented to both the database server and the backup server. Output from `ioscan` should show the robot and tape devices. For example:

```
# ioscan -fnkC tape
```

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
tape	20	0/9/0/0.1.55.255.0.0.0	stape	CLAIMED	DEVICE	HP Ultrium 2-SCSI
		/dev/rmt/20m	/dev/rmt/20mnb		/dev/rmt/c35t0d0BESTn	
		/dev/rmt/20mb	/dev/rmt/c35t0d0BEST		/dev/rmt/c35t0d0BESTnb	
		/dev/rmt/20mn	/dev/rmt/c35t0d0BESTb			
tape	21	0/9/0/0.1.55.255.0.0.1	stape	CLAIMED	DEVICE	HP Ultrium 2-SCSI
		/dev/rmt/21m	/dev/rmt/21mnb		/dev/rmt/c35t0d1BESTn	
		/dev/rmt/21mb	/dev/rmt/c35t0d1BEST		/dev/rmt/c35t0d1BESTnb	
		/dev/rmt/21mn	/dev/rmt/c35t0d1BESTb			
tape	22	0/9/0/0.2.55.255.0.0.0	stape	CLAIMED	DEVICE	HP Ultrium 2-SCSI
		/dev/rmt/22m	/dev/rmt/22mnb		/dev/rmt/c54t0d0BESTn	
		/dev/rmt/22mb	/dev/rmt/c54t0d0BEST		/dev/rmt/c54t0d0BESTnb	
		/dev/rmt/22mn	/dev/rmt/c54t0d0BESTb			
tape	23	0/9/0/0.2.55.255.0.0.1	stape	CLAIMED	DEVICE	HP Ultrium 2-SCSI
		/dev/rmt/23m	/dev/rmt/23mnb		/dev/rmt/c54t0d1BESTn	
		/dev/rmt/23mb	/dev/rmt/c54t0d1BEST		/dev/rmt/c54t0d1BESTnb	
		/dev/rmt/23mn	/dev/rmt/c54t0d1BESTb			
tape	14	0/10/0/0.1.39.255.0.0.0	stape	CLAIMED	DEVICE	HP Ultrium 2-SCSI
		/dev/rmt/14m	/dev/rmt/14mnb		/dev/rmt/c17t0d0BESTn	
		/dev/rmt/14mb	/dev/rmt/c17t0d0BEST		/dev/rmt/c17t0d0BESTnb	
		/dev/rmt/14mn	/dev/rmt/c17t0d0BESTb			
tape	15	0/10/0/0.1.39.255.0.0.1	stape	CLAIMED	DEVICE	HP Ultrium 2-SCSI
		/dev/rmt/15m	/dev/rmt/15mnb		/dev/rmt/c17t0d1BESTn	
		/dev/rmt/15mb	/dev/rmt/c17t0d1BEST		/dev/rmt/c17t0d1BESTnb	
		/dev/rmt/15mn	/dev/rmt/c17t0d1BESTb			
tape	16	0/10/0/0.2.39.255.0.0.1	stape	CLAIMED	DEVICE	HP Ultrium 2-SCSI
		/dev/rmt/16m	/dev/rmt/16mnb		/dev/rmt/c20t0d1BESTn	
		/dev/rmt/16mb	/dev/rmt/c20t0d1BEST		/dev/rmt/c20t0d1BESTnb	
		/dev/rmt/16mn	/dev/rmt/c20t0d1BESTb			
tape	17	0/10/0/0.2.39.255.0.0.2	stape	CLAIMED	DEVICE	HP Ultrium 2-SCSI
		/dev/rmt/17m	/dev/rmt/17mnb		/dev/rmt/c20t0d2BESTn	
		/dev/rmt/17mb	/dev/rmt/c20t0d2BEST		/dev/rmt/c20t0d2BESTnb	
		/dev/rmt/17mn	/dev/rmt/c20t0d2BESTb			

```
# ioscscan -fnkC autoch
```

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
autoch	2	0/10/0/0.2.39.255.0.0.0	schgr	CLAIMED	DEVICE	HP ESL9000 Series
		/dev/rac/c20t0d0				

After the tape library robot and tape devices have been presented to both the database server and the backup server then NetBackup can be installed.

Installing NetBackup components

The following NetBackup components are required for this solution:

- NetBackup Media Server
- NetBackup Oracle Database Agent and license
- NetBackup Shared Storage Option license

When running the NetBackup installation utility the required components and license keys must be installed. Please refer to the *VERITAS NetBackup 5.1 Installation Guide* for UNIX for instructions on installing NetBackup on HP-UX.

After the robotic and tape devices have been configured with the appropriate pass-through driver, the NetBackup Device Configuration Wizard should be run to properly configure the tape library for use by NetBackup. This should be done per the instructions in the HP-UX section of the *VERITAS NetBackup 5.1 Media Manager Device Configuration Guide* for UNIX and Windows.

The referenced documents can be found on the VERITAS NetBackup installation media or on the VERITAS web site at:

<http://support.veritas.com>

Creating NetBackup policies


The next step is to create the NetBackup policies. The following policies are required for this solution:

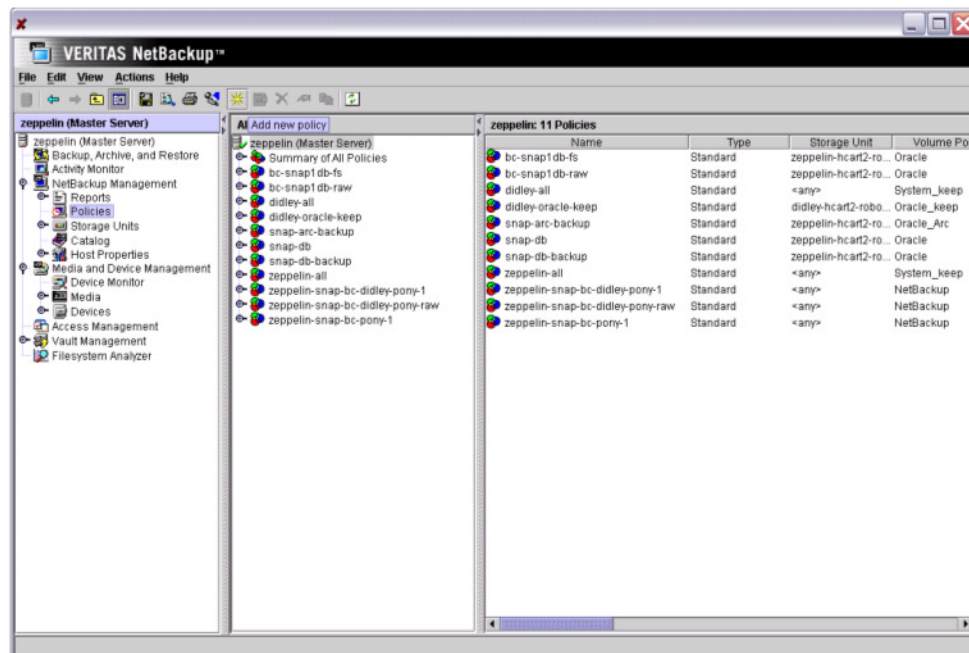
- **Controlling policy**—This policy is run by the database server, and controls the flow of the entire database backup process.
- **Database data file backup policy**—This policy is run on the backup server, and is not allowed to run until the controlling policy has completed its preprocessing script. There may be more than one database data file backup policy.
- **Archive log backup policy**—This policy is run on the backup server, and is started by the controlling policy postprocessing script.

After the policies have been created they should be deactivated until creation of the pre and post processing scripts has been completed. A policy can be deactivated by right clicking on the policy in the NetBackup Administration Console navigation pane and then clicking **Deactivate**.

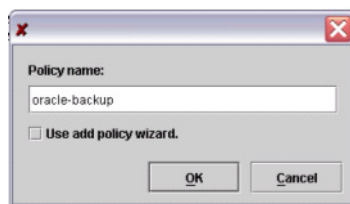
It is good practice to include the Oracle SID in the names of the NetBackup policies.

Creating the ZDB controlling policy

1. Start the NetBackup Administration Console on the NetBackup master server (this example uses a UNIX console).
 - If your master is a UNIX server, enter `/usr/opensv/netbackup/bin/jnbSA`
 - If your master is a Windows server, click **Start > All Programs > VERITAS NetBackup > NetBackup Administration Console**.
2. Click **Policies** under *NetBackup Management* in the navigation pane, and then click the starburst icon  in the toolbar to create a policy.

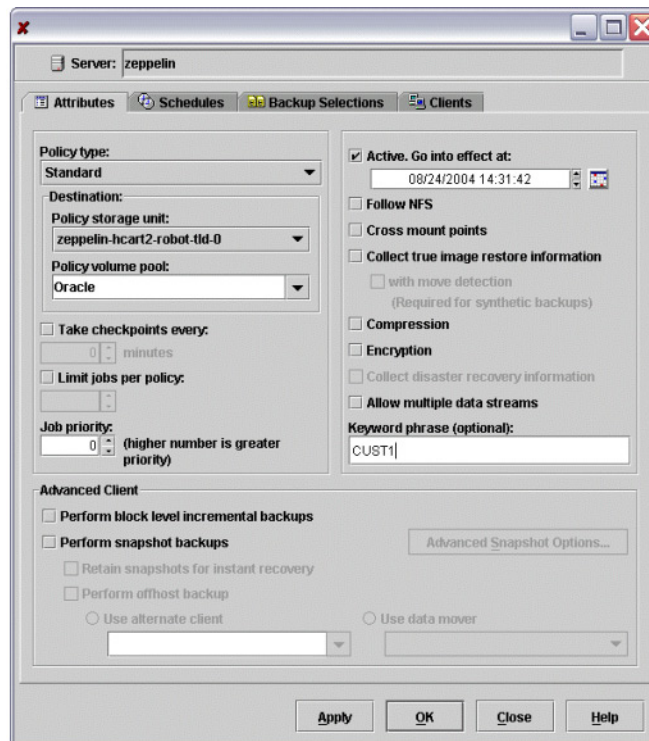


3. Enter the new policy name and click **OK** to start the Policy Create/Change window (*instance-backup* for example).

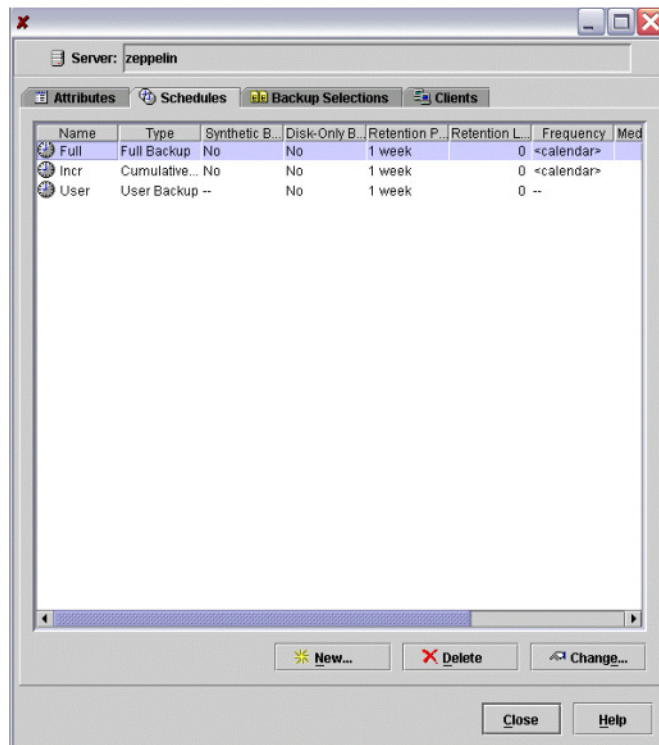


4. In the policy create window, under the Attributes tab, enter the following policy attributes:
 - a. **Policy type**—Standard
 - b. **Policy storage unit**—The backup server storage unit.
 - c. **Policy volume pool**—The media volume pool the policy should use. It is a good idea to create an Oracle volume pool to reserve media specifically for the database backups.

- d. **Keyword phrase**—A unique keyword must be set for all policies that are part of the Oracle ZDB solution. NetBackup uses the keyword to determine if all of the related policies are running during the database backup. The database instance name is a good keyword.



5. In the policy create window, under the Schedules tab create the following schedules (the frequency is dependent on your particular needs):
- Full**—Schedule type *Full Backup* required to automatically run full backups of the database.
 - Incr**—Schedule type *Cumulative Incremental Backup* to automatically run cumulative backups. This schedule is optional depending on your site requirements.
 - User**—Schedule type *User Backup* required for user initiated backups.



6. In the policy create window, under the *Backup Selections* tab, add the following files to back up (these are added here so that all files required for full database recovery are available on backup in case of disaster):
 - *Oracle instance parameter file*
 - *Oracle password file*
 - Additional files such as *.ora* files and control script to be included in the backup. The list of the additional files is site dependent.
7. In the policy create window, under the *Clients* tab add the database server as a client. Be sure to select the correct operating system type.
8. Click **Close** to complete the policy creation.

Creating the ZDB related data file backup policy

This policy is created in the same manner as the controlling policy. Multiple data file backup policies may be created for databases with many small or a few large data files. Creating multiple policies or setting up NetBackup to backup multiple streams in parallel may improve backup performance.

The steps are listed below but only the differences in this policy are pointed out:

1. Start the NetBackup Administration Console and navigate to the *Policies* window.
2. Click the starburst icon in the toolbar to create a new policy.
3. Enter the policy name (*instance-datafile-backup* for example).
4. Enter the policy attributes—The attributes should be the same as those entered for the controlling policy. This policy must have the same keyword.

5. Create the policy schedules—The schedules should be the same as those entered for the controlling policy. The start times should also be the same to ensure that both policies will run at the same time.
6. Enter the files that this policy should back up. This policy will be used to backup the database data files snapshot that is mounted on the backup server. Enter the snapshot mount point(s) or a complete path to each database data file snapshot.
7. Enter the policy client. The client must be the backup server.
8. Click **Close** to complete the policy creation.

Creating the ZDB archive log backup policy

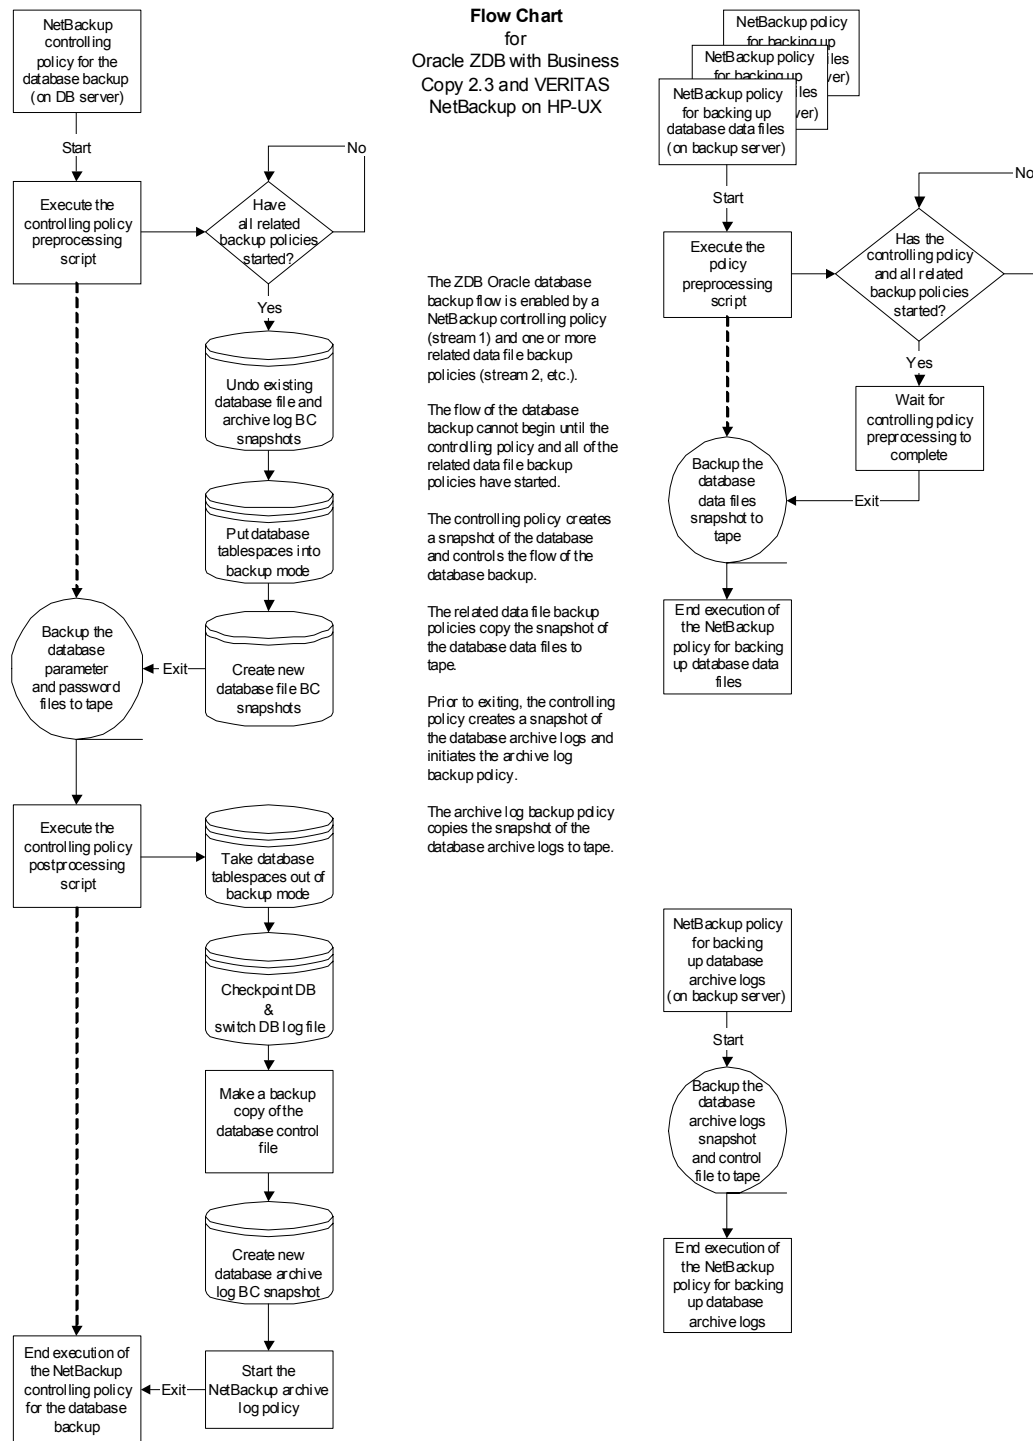
This policy is also created in the same manner as the controlling policy. The steps are listed below but only the differences in this policy are pointed out:

1. Start the NetBackup Administration Console and navigate to the *Policies* window.
2. Click the starburst icon in the toolbar to create a new policy.
3. Enter the policy name (*instance-arclog-backup* for example).
4. Enter the policy attributes—The policy type and storage unit should be the same as the controlling policy. The volume pool should be a special pool created specifically for archive log backups. The keyword must be different than the controlling policy keyword (*instance-arc* for example).
5. Create the policy schedules—The only required schedule is the *User Backup* schedule. There should not be a schedule created to automatically start backups because this policy is initiated by the controlling policy.
6. Enter the files that this policy should back up. This policy will be used to backup the archive log snapshot that is mounted on the backup server. Enter the snapshot mount point or a complete path to each archive log snapshot.
7. Enter the policy client. The client must be the backup server.
8. Click **Close** to complete the policy creation.

Note: It may be desirable to have the archive log backup policy remove archive logs that meet certain criteria, such as having been copied to tape. An archive log backup policy post processing script can be created to accomplish this task. Refer to the “NetBackup notify scripts” section of the *VERITAS NetBackup 5.1 System Administrator’s Guide, Volume II for UNIX* for details on creating post processing scripts.

Execution and flow of the NetBackup policy pre and post processing scripts

The following flow chart depicts the flow of the Oracle database backup as initiated by the related NetBackup policies created for the ZDB solution.



Creating the NetBackup policy pre and post processing scripts

The VERITAS NetBackup Oracle Database Agent provides the `setup_bli_scripts` utility for creating pre and post processing scripts. This utility is used for the Oracle ZDB solution to create the base scripts. The base scripts are then modified for the specific needs of the Oracle ZDB solution.

Note: Refer to the *VERITAS NetBackup 5.1 for Oracle System Administrator's Guide for UNIX* for details related to the Oracle Database Agent, details for creating script-based BLI backups, and detailed instructions for using the `setup_bli_scripts` utility.

Note: VERITAS supports the pre and post processing scripts that are created by the `setup_bli_scripts`. The modifications to the scripts for this solution are examples of how to modify your site specific pre and post processing scripts.

Creating scripts with the VERITAS NetBackup Oracle Database Agent Script Utility

The `setup_bli_scripts` utility prompts the user for input concerning the database and database backups. Prior to running the script, be prepared with the following information:

- The Oracle database administrator user name.
- The ORACLE_BASE directory path.
- The ORACLE_HOME directory path.
- The connect statement used to connect to the database.
- The Oracle database SID.
- The path to the Oracle instance init file.
- The path to the Oracle database config file if applicable.
- The path (including file name) where you would like NetBackup to make a backup copy of the control file. If this backup copy of the control file is placed on the archive log filesystem it can then be copied to tape as part of the archive log snapshot backup policy.

Note: It is also good practice to generate a control script based on database changes. The control script could be backed up as part of the controlling policy additional files.

- The path to the Oracle archive logs.
- Email address for output from the pre and post processing scripts if desired.
- The database backup method (the available methods allow for cold backups, hot backups, checkpoint, etc.)
- The list of policies included in the database backup. This list will include the controlling policy and the data file backup policies. This list should not include the archive log backup policy. The controlling policy is listed first.

Use the following steps to create the base pre and post processing scripts:

1. Start the script creation utility on the database server by entering,
`/usr/opensv/netbackup/ext/db_ext/oracle/bin/setup_bli_scripts`
2. Provide the input requested by the `setup_bli_scripts` utility. Carefully enter the information that you have gathered.

Modifying the ZDB policy preprocessing script

The preprocessing scripts `bpstart_notify.policy_name` were created by `setup_bli_scripts` utility for the controlling policy and the data file backup policy. The scripts are identical. This section outlines modifications required for the preprocessing scripts for the policy in control and the data file backup policy. After modifying the data file backup policy preprocessing script it should be copied to the backup server `/usr/opensv/netbackup/bin` directory.

Items to modify in the policy in control preprocessing script:

- **KEYWORD and STREAMS parameters**—These parameters will not be passed to the preprocessing script because the *Perform block level incremental backups* option was not selected while creating the policy in control. These parameters must be set manually in the preprocessing script. The **KEYWORD** parameter should be set to the keyword used in the policy in control. The **STREAMS** parameter should be set to the number of policies with the same keyword. There are at least 2 policies with the same keyword, the policy in control and at least one data file backup policy. These parameters should be added to the parameter initialization section similar to the following (in this example SNAP2 is the keyword used while creating the policies):

```
# -----  
# Assign names to our parameters.  
#-----  
  
CLIENTNAME=$1  
POLICYNAME=$2  
SCHEDNAME=$3  
SCHEDTYPE=$4  
KEYWORD= " SNAP2 "  
STREAMS=2
```

- **CONTROL_DIR parameter**—This parameter contains the path of the directory where all preprocessing logs will be written. This does not have to be modified, however by default it logs to a subdirectory of the NetBackup bin directory. This can be modified to be a subdirectory of the NetBackup log directory.
- **START_FILE and START_LINK parameters**—These parameters contain the paths and filenames for the controlling policy start file and start link. The controlling policy preprocessing script creates the start file at startup. All other policies with the same policy keyword create a start link at startup. Once the start file and all related policy start links have been created, the preprocessing scripts are allowed to continue. For the Oracle ZDB solution this is a problem because the policies do not all run on the same server. To resolve this issue the **START_FILE** and **START_LINK** parameters are changed to a path on an NFS file system that is shared among the database server and backup server.

- **Remove or comment the block incremental check**—This check causes the preprocessing script to exit if the backup is not a block incremental. The `setup_bli_scripts` utility is used to create pre and post processing scripts for Oracle block incremental backups. The Oracle ZDB solution is not block incremental so this section should be removed or commented. The section of the script to be removed or commented should be similar to the following:

```
if [ "${SCHEDTYPE}" = "UBAK" -o $STREAMS -eq 0 ]
then
    # If this is not a block incremental, exit now.
    exit 0
fi
```

- Modify the START_LINK creation section to create a file instead of a hard link—This section of code must be modified to create a START_LINK file rather than a hard link. This is required because a hard link cannot be created on an NFS mounted file system. The section of the script to be modified should be similar to the following (changes are in bold):

```
# -----
# Create a hard link to the START_FILE. We will use the link count to
# decide when all the streams have been started.
# -----

if [ -f $START_FILE ]
then
#     /bin/ln $START_FILE $START_LINK
#     LINKCOUNT=`ls -l $START_FILE|sed 's/ / /g'|sed 's/ / /g'|cut -f2 -d" "`
#     /bin/touch $START_LINK
#     /bin/chmod 777 $START_LINK
#     LINKCOUNT=`ls ${START_FILE}*|wc -w`
else
    LINKCOUNT=0
fi

while [ $STREAMS -ge $LINKCOUNT ]
do
    if [ -f "${ERROR_FILE}" ]
    then
        $ECHO "`$TIME` $POLICYNAME Error file $ERROR_FILE detected" >>$OUTF
        $ECHO "`$TIME` $POLICYNAME while waiting for link count of \
$LINKCOUNT on $START_FILE file" >>$OUTF
        abort
    fi

    /bin/sleep 1

    if [ -f $START_FILE ]
    then
        # The "policy in control" may remove our link if
        # the "policy in control" isn't the first one started.
#         /bin/ln $START_FILE $START_LINK
#         /bin/touch $START_LINK
#         /bin/chmod 777 $START_LINK

#         LINKCOUNT=`ls -l $START_FILE|sed 's/ / /g'|sed 's/ / /g'|cut -f2\
#         -d" "`
#         LINKCOUNT=`ls ${START_FILE}*|wc -w`
    else
        LINKCOUNT=0
    fi
done
```

- **Add a section to undo existing data file archive log snapshots**—This section of code is required for the controlling policy to undo snapshots prior to placing the tablespaces in backup mode. The following sample section of the preprocessing script includes the undo of the snapshots in bold text:

```
# -----
# All streams have been started, so we can now continue.
# -----

$ECHO "" >>$OUTF
$ECHO "`$TIME` $POLICYNAME All streams have been started." >>$OUTF
$ECHO "" >>$OUTF

if [ "$POLICYNAME" = "$POLICY_IN_CONTROL" ]
then
    /bin/rm -f $SHUTDOWN_BKUP_RESTART_FILE
    /bin/rm -f $SHUTDOWN_CKPT_RESTART_FILE
    /bin/rm -f $ALTER_TABLESPACE_FILE
    /bin/rm -f $NODATA_CKPT_HOT_FILE
    /bin/rm -f "${POLICY_IN_CONTROL_FILE}"

    # Write the POLICY_IN_CONTROL value to a file so that bpbkar can read it.

    /bin/touch $POLICY_IN_CONTROL_FILE
    /bin/chmod 777 $POLICY_IN_CONTROL_FILE
    $ECHO $POLICY_IN_CONTROL >>"${POLICY_IN_CONTROL_FILE}"

    #
    # Section added by HP to undo current snaps
    #
    $ECHO "Undoing existing DB volume snaps." >>$OUTF
    /home/bc/bc-snap -u -a sanapp01 bc-oracledb bc-oraclelog bc-oraclearc >>$OUTF

    # Find out if the instance is up.

    SMONPROC=`${PSCMD}|grep "smon_${ORACLE_SID}"|grep -v grep|wc -l`

    if [ ${SMONPROC} -eq 0 ]
    then
        $ECHO "`$TIME` $POLICYNAME Database ${ORACLE_SID} is down." >>$OUTF
    fi

    # Shut down the DB or put it in "backup" mode.
```

The `bc-snap` script is the script that was created to run and undo BC snapshot jobs. The `-u` option tells the script to undo the jobs. The `-a sanapp01` option is the Storage Management Appliance. The `bc-oracledb bc-oraclelog bc-oraclearc` parameters are the list of jobs to undo.

- **Add a section to create new data file snapshots**—This section of code is required for the controlling policy to create snapshots after placing the tablespaces in backup mode. The following sample section of the preprocessing script includes the creation of the snapshots in bold text (this section of code follows immediately after the section above where the snapshots were undone):

```
# Shut down the DB or put it in "backup" mode.

case "$METHOD"
in
  "SHUTDOWN_BKUP_RESTART")
    ... ..
    ... ..
    ;;
  "SHUTDOWN_CKPT_RESTART")
    ... ..
    ... ..
    ;;
  "ALTER_TABLESPACE")
    ... ..
    ... ..
    ;;
  "NODATA_CKPT_HOT")
    ... ..
    ... ..
    ;;
esac

#
# Section added by HP to do new snaps
#
$ECHO "Creating new DB volume snaps." >>$OUTF
/home/bc/bc-snap -d -a sanapp01 bc-oracledb bc-oraclelog >>$OUTF

# We now have all the streams started. Now wait for others to unlink.

if [ -f $START_FILE ]
then
  LINKCOUNT=`ls -l $START_FILE|sed 's/ / /g'|sed 's/ / /g'|cut -f2 -d" "`
else
  LINKCOUNT=0
fi
```

The `bc-snap` script is used in this case to create snapshots. The `-d` option tells the script to create new snapshots. The `-a sanapp01` option is the Storage Management Appliance. The `bc-oracledb bc-oraclelog` parameters are the list of jobs to run. The `bc-oraclearc` job is not run at this time but will be run in the postprocessing script.

Items to modify in the data file backup policy preprocessing script:

The data file backup policy preprocessing script should have the same modifications as the policy in control preprocessing script except for the following items:

- **Remove or comment the check for the existence of the \$ORACLE_INIT file**—The data file backup policy runs on the backup server not the database server. The \$ORACLE_INIT file will not exist on the backup server unless the backup server is used to verify the database snapshots by starting the database on the backup server. The following section should be removed from the preprocessing script for the data file backup policy:

```
if [ ! -f ${ORACLE_INIT} ]
then
    $ECHO "`$TIME` $POLICYNAME ORACLE_INIT ${ORACLE_INIT} not found" >>$OUTF
    abort
fi
```

- **Do not add the sections to undo and create the data file snapshots**—The data file snapshots are undone and created by the policy in control preprocessing script and are not required in the data file backup preprocessing script.

Modifying the ZDB policy postprocessing script

The postprocessing scripts `bpend_notify.policy_name` were created by `setup_bli_scripts` utility for the controlling policy and the data file backup policy. The scripts are identical. This section outlines modifications required for the postprocessing scripts for the policy in control and the data file backup policy. After modifying the data file backup policy postprocessing script it should be copied to the backup server `/usr/opensv/netbackup/bin` directory.

Items to modify in the policy in control postprocessing script:

- **KEYWORD and STREAMS parameters**—These parameters will not be passed to the postprocessing script because the *Perform block level incremental backups* option was not selected while creating the policy in control. These parameters must be set manually in the postprocessing script. The KEYWORD parameter should be set to the keyword used in the policy in control. The STREAMS parameter should be set to the number of policies with the same keyword. There are at least 2 policies with the same keyword, the policy in control and at least one data file backup policy. These parameters should be added to the parameter initialization section similar to the following (in this example SNAP2 is the keyword used while creating the policies):

```
# -----
# Assign names to our parameters.
# -----

CLIENTNAME=$1
POLICYNAME=$2
SCHEDNAME=$3
SCHEDTYPE=$4
KEYWORD= " SNAP2 "
STREAMS=2
```

- **CONTROL_DIR parameter**—This parameter contains the path of the directory where all postprocessing logs will be written. This does not have to be modified, however by default it logs to a subdirectory of the NetBackup bin directory. This can be modified to be a subdirectory of the NetBackup log directory.
- **START_FILE parameter**—This parameter contains the path and filename for the controlling policy start file. The controlling policy preprocessing script creates this file at startup. The postprocessing script removes this file if it was left by the preprocessing script. This parameter should be set to the same value as in the preprocessing script.
- **Remove or comment the block incremental check**—This check causes the postprocessing script to exit if the backup is not a block incremental. The `setup_bli_scripts` utility is used to create pre and post processing scripts for Oracle block incremental backups. The Oracle ZDB solution is not block incremental so this section should be removed or commented. The section of the script to be removed or commented should be similar to the following:

```
if [ "${SCHEDTYPE}" = "UBAK" -o $STREAMS -eq 0 ]
then
    # If this is not a block incremental, exit now.
    exit 0
fi
```

- **Remove or comment the section that creates the END_FILE**—This section is not required for the Oracle ZDB solution. The policy in control and the data file backup policy do not need to be in sync after the data volume snapshot backups have started. The section of the script to be removed or commented should be similar to the following (the commented section is in bold):

```
#if [ "$POLICYNAME" = "$POLICY_IN_CONTROL" ]
#then
#    if [ -f $END_FILE ]
#    then
#        # If the control file already exists, remove it along with
#        # the link files from all the streams.
#        /bin/rm -f ${END_FILE}* ${END_LINK}*
#    fi
#    /bin/touch $END_FILE
#    /bin/chmod 777 $END_FILE
#else
#    /bin/rm -f $END_LINK

# Wait for the control file to be created by the "policy in control"
# before continuing.

#    SLEEP_COUNT=0

#    while [ ! -f $END_FILE ]
#    do
#        /bin/sleep 1
#        SLEEP_COUNT=`expr $SLEEP_COUNT + 1`

#        if [ $SLEEP_COUNT -gt 7200 ]
#        then
#            # If we hit 2 hours, create the error file ourselves
#            # so that we don't wait forever.

#            $ECHO "`$TIME` $POLICYNAME Exceeded 7200 second wait count." /
#>>$OUTF
#            $ECHO "`$TIME` $POLICYNAME Creating error file $ERROR_FILE" /
#>>$OUTF
#            /bin/touch "${ERROR_FILE}"
#            /bin/chmod 777 "${ERROR_FILE}"
#            abort
#        elif [ $SLEEP_COUNT -gt $BPEND_TIMEOUT ]
#        then
#            # If we've exceeded the timeout and the policy in
#            # control hasn't started yet, start looking for the
#            # error file.

#            if [ -f "${ERROR_FILE}" ]
#            then
#                $ECHO "`$TIME` $POLICYNAME Error file $ERROR_FILE /
#detected" >>$OUTF
```

```
#                                     $ECHO "`$TIME` $POLICYNAME while waiting for $END_FILE /
#file" >>$OUTF
#                                     abort
#                                     fi
#                                     fi
#     done
#fi
```

```
# -----
# Create a hard link to the END_FILE. We will use the link count to
# decide when all the streams have been started.
# -----
```

- **Remove or comment the archive log backup section**—This section initiates backups of the archive logs and control file by the database server. This should be removed or commented because the Oracle ZDB solution handles backup of the archive logs and control file via a snapshot on the backup server. The section of the script to be removed or commented should be similar to the following (the commented section is in bold):

```
# -----
# If everything is OK, backup the archive logs and control file.
# -----

if [ ! -f "${ERROR_FILE}" -a $STATUS -eq 0 ]
then
    if [ "$POLICYNAME" = "$POLICY_IN_CONTROL" ]
    then
        # Checkpoint the database and switch archive log files.

        checkpoint_database

        # Initiate a user directed backup of the control file and
        # the archive log files.

        if [ "$ORACLE_LOGS" = "xxxxxx" ]
        then
            # The ORACLE_LOGS was not specified.

            $BKUP -k "${KEYWORD}" -c $POLICY_IN_CONTROL $ORACLE_INIT \
            $ORACLE_CONFIG $ORACLE_CNTRL

        else
            $BKUP -k "${KEYWORD}" -c $POLICY_IN_CONTROL $ORACLE_INIT \
            $ORACLE_CONFIG $ORACLE_CNTRL $ORACLE_LOGS

        fi

        # Snap archive logs and backup
        /home/bc/bc-snap -d -a sanapp01 bc-oraclearc >>$OUTF
        $BKUP -p snap-arc-backup -i -h zeppelin -s Full

    fi
fi
```

- **Add a section to create new archive log snapshots and start the archive log backup policy**—This section of code is required for the controlling policy to create the archive log snapshot after placing the tablespaces in normal mode, doing a database checkpoint and switching the log file. It is also required to start the archive log backup policy.

The following sample section of the postprocessing script includes the creation of the snapshot and initiation of the archive log backup policy in bold text (it is the same section as the previous modification where the control file and archive log backups were removed or commented):

```
# -----
# If everything is OK, backup the archive logs and control file.
# -----

if [ ! -f "${ERROR_FILE}" -a $STATUS -eq 0 ]
then
    if [ "$POLICYNAME" = "$POLICY_IN_CONTROL" ]
    then
        # Checkpoint the database and switch archive log files.

        checkpoint_database

        # Initiate a user directed backup of the control file and
        # the archive log files.

        if [ "$ORACLE_LOGS" = "xxxxxx" ]
        then
            # The ORACLE_LOGS was not specified.
            $BKUP -k "${KEYWORD}" -c $POLICY_IN_CONTROL $ORACLE_INIT \
            $ORACLE_CONFIG $ORACLE_CNTRL
        else
            $BKUP -k "${KEYWORD}" -c $POLICY_IN_CONTROL $ORACLE_INIT \
            $ORACLE_CONFIG $ORACLE_CNTRL $ORACLE_LOGS
        fi

        # Snap archive logs and backup
        /home/bc/bc-snap -d -a sanapp01 bc-oraclearc >>$OUTF
        $BKUP -p snap-arc-backup -i -h zeppelin -s Full
    fi
fi
```

The `bc-snap` script is used to create a snapshot of the archive logs. The `-d` option tells the script to create the snapshot. The `-a sanapp01` options is the name of the Storage Management Appliance. The `bc-oraclearc` parameter is the list of jobs to run. If the backup copy of the control file was copied to the archive log file system on the backup server then the backup of the archive log snapshot will include the control file.

Items to modify in the data file backup policy postprocessing script:

The data file backup policy postprocessing script should have the same modifications as the policy in control postprocessing script except for the following items:

- **Remove or comment the check for the existence of the \$ORACLE_INIT file**—The data file backup policy runs on the backup server not the database server. The \$ORACLE_INIT file will not exist on the backup server unless the backup server is used to verify the database snapshots by starting the database on the backup server. The following section should be removed from the postprocessing script for the data file backup policy:

```
if [ ! -f ${ORACLE_INIT} ]
then
    $ECHO "`$TIME` $POLICYNAME ORACLE_INIT ${ORACLE_INIT} not found" >>$OUTF
    abort
fi
```

- **Do not add the section to create the archive log snapshots and start the archive log backup**—The archive log snapshots creation and backup initiation are done by the policy in control postprocessing script and are not required in the data file backup postprocessing script.

Running the NetBackup policies to accomplish a ZDB of the Oracle database

It is time to do a ZDB backup of the database. The policies and scripts required have all been created. The backup can be initiated manually or through the policy schedules. It is best to test the backup manually prior to automating through the policy schedules.

Running the backup manually

The database backup can be activated manually using the following steps:

1. **Activate the policies**—The policies created for the backup can be activated by right clicking on the policy in the NetBackup Administration Console navigation pane and then clicking **Activate**.
2. **Start the controlling policy**—The controlling policy can be manually started by right clicking on the policy in the NetBackup Administration Console navigation pane and then clicking **Manual Backup**.
3. **Start the data file backup policy**—The data file backup policy can be manually started by right clicking on the policy in the NetBackup Administration Console navigation pane and then clicking **Manual Backup**.

To follow the backup progress, go to the NetBackup Administration Console Activity monitor. Note that the policies will have a status of *connecting* while the preprocessing scripts are running.

The pre and post processing logs, created by the scripts while running, are text files that can be viewed to see the activity of the scripts. The logs are located in the *CONTROL_DIR* directory specified in the scripts.

The Business Copy Web GUI can be used to view the progress of BC jobs.

Automating the backup

The backup can be automated by using the NetBackup policy schedules that were created during the controlling policy and data file backup policy creation. Using the schedules, backups can be setup to run at a specific time on a nightly, weekly, monthly, etc. basis. The schedules for the controlling policy and the data file backup policy should be identical. Both policies must run at the same time for the backup to succeed. The archive log backup policy should not be set to run automatically because it is initiated by the controlling policy postprocessing script.

Verifying database integrity with snapshot volumes

For added security and comfort with the Oracle ZDB solution, it may be desirable to verify Oracle database integrity using the snapshot volumes. This can be done by starting the database instance on the backup server using the data on the snapshot volumes.

Requirements for verifying the database on the backup server using the snapshot volumes:

- Oracle must be installed on the backup server.
- Current Oracle data file, redo log, and archive log snapshot volumes must be mounted on the backup server using the same mount point the primary data volumes use on the database server.
- The control file, password file, and parameter file must be copied or restored to the backup server. (There may also be other site specific *.ora* files needed).

If the requirements are met, the database instance can be recovered (not open) on the backup server by using normal Oracle database recovery procedures. After the database has been recovered it can be started (read only) and queries can be run against the database to verify the database integrity.

Oracle Database Restore and Recovery

3

Database recovery overview

The process of recovering the database consists of:

1. **Identifying and preparing resources**—Prior to recovery, you must know where you will be recovering from. Is the data on a snapshot volume or on tape? What resources are needed to recover from a snapshot or from tape?
2. **Preparing the database for recovery**—The database must be shutdown in NORMAL, IMMEDIATE, or ABORT mode.
3. Restoring needed files:
 - a. All data files that belong to the tablespaces being restored.
 - b. Any archived redo logs needed to make restored files current.
 - c. One copy of the control file, if it was damaged.
 - d. The parameter file, if different than the current the file.
 - e. The *orapwd* file, if it was damaged.
4. **Recovering the database and applying redo logs with sqlplus**—This step depends on the nature of the failure and whether the database needs to be online during recovery. Oracle provides methods for performing data file recovery, tablespace recovery, and database recovery. Consult the Oracle documentation for more information on these functions.

Restoring the database files from the snapshot area

If the snapshot that contains the necessary backup set is available, the database can be restored directly from the snapshot, preventing the need to retrieve the data from tape.

If the snapshot is mounted on the backup server, it can be unmounted, unrepresented, then re-presented to the database server. Use the HP HSV Element Manager to unrepresent and re-present EVA volumes. The snapshot can be made visible to the OS by scanning for new devices using *ioscan*.

Note: If the snapshot volume on the backup server is an LVM volume then the volume group will need to be exported on the backup server and imported on the database server prior to mounting the volume on the database server.

After the database server has access to the snapshot, the necessary files can be copied from the snapshot back to the source volumes.

Note: The snapshot volume should be unmounted and unrepresented to the database server, and re-presented and mounted on the backup server following the restore so the BC job can successfully undo the snapshot and recreate it for the next backup.

If the snapshot remains mounted on the backup server then remote copy or ftp can be used to copy the necessary files back to the database server source volumes.

Note: When copying data files from the snapshot mounted on the backup server, create an archive of the data files (e.g. *tar* or *cpio*), copy the archive to the database server, and extract the data files to the data file volume.

Restoring the database files from tape

Restoring the database backup created by the ZDB solution from tape is no different than a normal NetBackup restore. If cumulative incremental backups were configured then restoring to one of the incremental backup images requires a restore of the last full backup image and the cumulative incremental backup image. You can start the restores from the NetBackup server by using the Backup, Archive, and Restore interface.

Refer to the *NetBackup User's Guide for UNIX* and the *NetBackup for Oracle System Administrator's Guide for UNIX* for more detailed information on restoring.

These documents can be found on the VERITAS NetBackup installation media or on the VERITAS web site at:

<http://support.veritas.com>

To restore database data files and archive logs directly to the database server the restore client must be set to the database server (the backup server was the client that performed the backup). The restore should be initiated from the NetBackup server using the Backup, Archive, and Restore interface.

Index

A

- array integration 14
- audience 5
- authorized reseller, HP 8

B

- backup
 - automating 53
 - running manually 53
- backup server, defined 18
- backup server, setting up 34
- BC Job, defined 18
- BC run, defined 18
- BC undo, defined 18
- Business Continuance Volume, defined 18
- Business Copy
 - creating snapshot jobs 25
 - creating snapshot script 30
 - installing 20
 - installing host agent 22
 - overview 11
 - setting up for the EVA 20
 - web GUI 11

C

- conventions
 - document 6
 - equipment symbols 7
 - text symbols 6

D

- database
 - availability 14
 - recovery overview 55
 - restoring files from snapshot 55
 - restoring files from tape 56
- server
 - defined 18
 - setting up 34
- disk groups, defined 10
- document
 - conventions 6
 - prerequisites 5
 - related documentation 6

E

- Enterprise Backup Solutions
 - overview 9
- equipment symbols 7

G

- getting help 8

H

- help, obtaining 8
- hosts, defined 10
- HP
 - authorized reseller 8
 - HSV Element Manager 10
 - storage web site 8
 - technical support 8
- HSV Element Manager 10
 - installing Business Copy on 20

L

LUNs, identifying on the database server [25](#)

N

NetBackup

- creating policies [35](#)
- creating policy processing scripts [41](#)
- flow of policy processing scripts [40](#)
- installing components [35](#)
- postprocessing script, defined [18](#)
- preprocessing script, defined [18](#)
- running policies [53](#)

P

prerequisites [5](#)

R

related documentation [6](#)

restores [14](#)

S

SAN zoning [15](#)

snapclone, defined [18](#)

snapshot jobs

- creating [25](#)
- creating script to automate [30](#)
- using CLI Utility [30](#)
- verifying on the EVA [28](#)

snapshot script, creating [30](#)

snapshot, defined [18](#)

solution components [17](#)

solution steps [18](#)

Storage Management Appliance [10](#)

symbols in text [6](#)

symbols on equipment [7](#)

T

technical support, HP [8](#)

text symbols [6](#)

V

virtual disk, defined [10](#), [18](#)

W

warning

symbols on equipment [7](#)

web sites

HP storage [8](#)

Z

zero downtime backup

creating archive log backup policy [39](#)

creating data file backup policy [38](#)

creating the controlling policy [36](#)

defined [9](#)

HP support of [9](#)

modifying the policy postprocessing script [47](#)

modifying the policy preprocessing script [42](#)

solution, designing [14](#)

topology [13](#)

zoning, SAN [15](#)